

RECEIVED
CENTRAL FAX CENTER

MAY 02 2007

IN THE CLAIMS

The pending unamended claims are reproduced below:

1. (PREVIOUSLY PRESENTED) A method of optimizing execution of a query that accesses data stored on a data store connected to a computer, comprising:
 using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query; and
 executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.
2. (ORIGINAL) The method of claim 1, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.
3. (ORIGINAL) The method of claim 1, further comprising:
 generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and
 using the generated cardinality estimates to determine an optimal query execution plan for the query.
4. (ORIGINAL) The method of claim 3, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.
5. (ORIGINAL) The method of claim 4, wherein the predicates are applied by one or more of the pre-defined queries.
6. (ORIGINAL) The method of claim 5, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.
7. (ORIGINAL) The method of claim 4, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

8. (ORIGINAL) The method of claim 7, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

9. (ORIGINAL) The method of claim 8, further comprising determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

10. (ORIGINAL) The method of claim 9, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

11. (ORIGINAL) The method of claim 10, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

12. (PREVIOUSLY PRESENTED) An apparatus for optimizing execution of a query, comprising:

a computer having a data store coupled thereto, wherein the data store stores data;
one or more computer programs, performed by the computer, for using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query, and executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.

13. (ORIGINAL) The apparatus of claim 12, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.

14. (ORIGINAL) The apparatus of claim 12, further comprising:
one or more computer programs for generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and

one or more computer programs for using the generated cardinality estimates to determine an optimal query execution plan for the query.

15. (ORIGINAL) The apparatus of claim 14, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.

16. (ORIGINAL) The apparatus of claim 15, wherein the predicates are applied by one or more of the pre-defined queries.

17. (ORIGINAL) The apparatus of claim 16, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

18. (ORIGINAL) The apparatus of claim 15, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

19. (ORIGINAL) The apparatus of claim 18, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

20. (ORIGINAL) The apparatus of claim 19, further comprising one or more computer programs for determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

21. (ORIGINAL) The apparatus of claim 20, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

22. (ORIGINAL) The apparatus of claim 21, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

23. (PREVIOUSLY PRESENTED) An article of manufacture comprising a program storage medium readable by a computer and embodying one or more instructions executable by the computer to optimizing execution of a query that accesses data stored on a data store connected to the computer, comprising:

using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query; and

executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.

24. (ORIGINAL) The article of claim 23, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.

25. (ORIGINAL) The article of claim 23, further comprising:

generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and

using the generated cardinality estimates to determine an optimal query execution plan for the query.

26. (ORIGINAL) The article of claim 25, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.

27. (ORIGINAL) The article of claim 26, wherein the predicates are applied by one or more of the pre-defined queries.

28. (ORIGINAL) The article of claim 27, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

29. (ORIGINAL) The article of claim 26, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

30. (ORIGINAL) The article of claim 29, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

31. (ORIGINAL) The article of claim 30, further comprising determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

32. (ORIGINAL) The article of claim 31, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

33. (ORIGINAL) The article of claim 32, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.